arXiv:1704.00805v2 [math.OC] 9 Apr 2017

On the Properties of the Softmax Function with Application in Game Theory and Reinforcement Learning

Bolin Gao and Lacra Pavel

***Abstract*—In this paper, we utilize results from convex analysis and monotone operator theory to derive additional properties of the softmax function not yet covered in the existing literature. In particular, we show that the softmax function is the monotone gradient map of the log-sum-exp function. We show that the inverse temperature parameter λ determines the Lipschitz and co-coercivity properties of the softmax function. We demonstrate the usefulness of these properties through an application in population games and reinforcement learning.**

I. Introduction

The softmax function is one of the most well known functions in science and engineering and has enjoyed widespread usage in fields such as game theory [1], [2], reinforcement learning [3] and machine learning [4], [5]. The primary application of the softmax function is to map a vector in a Euclidean space to a probability distribution. From a game theory and reinforcement learning perspective, the softmax function maps the payoff from a game or score associated with a payoff to a mixed strategy [1], [2], [3]. From the perspective of logistic regression and feature classification, the softmax function maps a vector of logits or feature variables to a posterior probability distribution [4], [5]. The broader engineering applications involving the softmax function are numerous, some interesting examples can be found in the field of VLSI and neuromorphic computing, see [29], [30], [31], [33].

The term "softmax" is a portmanteau of "soft" and "argmax" [4] and the function first appeared in the work of Luce [9], although its coinage is mostly credited to Bridle [10]. Depending on the context in which the softmax function appears, it also goes by the name of Boltzmann distribution [1], [3], [28], Gibbs map [16], [40], logit map [1], [2], [17], [52], quantal response function [11], or smooth best response function [38]. In this paper, we will refer to the softmax function interchangeably as the softmax operator, softmax map, softmax choice or softmax selection rule. One should take care in distinguishing the softmax function used in the present context from the log-sum-exp function, which is often also referred to as the "softmax" because the log-sum-exp function is soft approximation of the vector-max function [6], [18].

There are many reasons behind the wide-spread usage of the softmax function. The most appealing feature of the softmax function is that it is a differentiable, smooth map. In the context of reinforcement learning, the softmax function ensures a trade-off between exploitation and exploration, in that every strategy in an agent's possession has a chance of being explored. Unlike some other choice mechanisms such as ϵ-greedy [3], the usage of softmax selection rule is favorably supported by experimental literature in game theory and reinforcement learning as a plausible model for realistic decision making. For instance, in [50], it was noted that the behavior of monkeys during reinforcement learning experiments is consistent with the softmax selection rule. From the perspective of game theory, the softmax function characterizes the so called "quantal response equilibrium" [11], which accounts for incomplete information and random perturbation of the payoff during game play and has been noted for its better versatility in describing the outcomes of game play as compared to the Nash equilibrium [1], [11]. For additional discussion on the connection between softmax selection rule and human decision making process, see [25], [26], [27], [35]. Finally it has been noted that the input-output behavior of the softmax function can be thought of as lateral inhibition in biological neural networks [4].

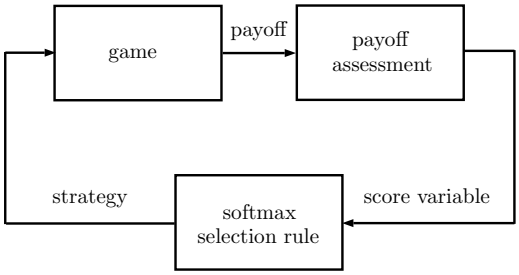

Fig. 1: High-level feedback system representation of a reinforcement learning scheme with softmax selection rule

Despite the intuition researchers have acquired over the years with respect to the usage of the softmax function, its mathematical properties are still far from being fully understood. However, it is apparent that there is the need to have a better understanding of these properties. For instance, in control-theoretic analysis of stateless multi-agent reinforcement learning schemes, the overall reinforcement learning scheme can be represented as an interconnected system as in Figure 1. When the action selection rule is taken as the softmax function, it is of interest which, if any, properties of softmax will allow us to conclude convergence of the learning algorithm towards a solution of the game (e.g. a Nash equilibrium). Although the properties of interest are fairly mundane

B. Gao and L. Pavel are with the Department of Electrical and Computer Engineering, University of Toronto, Toronto, ON, M5S 3G4, Canada. Emails: bolin.gao@mail.utoronto.ca, pavel@ece.utoronto.ca

from a control theory perspective, virtually no references to these properties can be found in the existing literature. With regard to applications in the context of reinforcement and machine learning, the adjustment of the temperature constant of the softmax function is still mostly performed on a rule-of-thumb basis. It has been speculated in [36] that the temperature constant can be used for reinforcement learning algorithms to get out of local minima. Therefore, an adaptive mechanism of adjusting the temperature constant would be desirable in many applications. Clearly, these questions can only be answered if we delve deeper into the properties of the softmax function.

The goal of this paper is to expand on the known mathematical properties of the softmax function. We will perform our analysis and derive new properties of this function using tools from convex analysis [6], [18] and monotone operator theory [19], [20]. Throughout the paper, we motivate some of the results through insights from game theory, particular from the field of evolutionary game theory [12], [13], [14] and population games [2]. This is done because it has been known that stateless multi-agent reinforcement learning with softmax choice has close connections with the field of evolutionary game theory [7], [8], [16], [17], [48], [50], [53]. In particular, evolutionary game theory has been proposed as the theoretic framework of multi-agent reinforcement learning [7], [8]. It is our hope that researchers could adapt our results presented here and apply them to their domain-specific problems.

The organization of this paper is as follows. Section II introduces notation convention for the rest of the paper. Section III introduces the definition of the softmax function, its different representations and some of its known properties from the existing literature. Section IV provides the background to convex optimization and monotone operator theory. In Section V, we derive additional properties of the softmax function. Section VI provides a control-theoretical analysis of a stateless continuous-time score-based reinforcement learning scheme within a population games setup to illustrate the usefulness of these additional properties of the softmax function. Section VII provides the conclusion and some open problems.

II. Notations

We denote the p-norm of a vector as $\|\cdot\|_p$, where $1 \leq p \leq \infty$, a vector in $\mathbb{R}^n$ as $z \in \mathbb{R}^n$, and a vector in the $n-1$ unit simplex given by $\Delta^{n-1} = \{x \in \mathbb{R}^n | \|x\|_1 = 1, x_i \geq 0\}$ as $x \in \Delta^{n-1}$. We denote the interior of Δ^{n-1} as $\operatorname{int}(\Delta^{n-1}) = \{x \in \mathbb{R}^n | \|x\|_1 = 1, x_i > 0\}$. Let $e_i \in \mathbb{R}^n$ denote the i^{th} canonical basis of $\mathbb{R}^n$, e.g. $e_i = \begin{bmatrix}0, \dots, 1, \dots, 0\end{bmatrix}^T$, where 1 occupies the i^{th} position, the vector of ones as $\mathbf{1} = \begin{bmatrix}1, \dots, 1\end{bmatrix}^T$ and the vector of zeros as $\mathbf{0} = \begin{bmatrix}0, \dots, 0\end{bmatrix}^T$. Finally, we denote matrices using bold capital letters such as $\mathbf{A}$.

III. Review of the Softmax Function and Its Basic Properties

While the *softmax* function may take on different appearances depending on the context of its application, the base model is that of a vector-valued function, whose individual components consists of an exponential evaluated at an element of a vector normalized by the sum of the exponential of all elements of that vector. In this paper we choose to study the most well known and well accepted version of the softmax function [4], [31], [34], [35], [37]. In the following, we present several well-known and equivalent representations of the softmax function and explore some of its properties that are either immediate or have been covered in existing literatures.

Definition 1. The softmax function is given by $\sigma : \mathbb{R}^n \to \operatorname{int}(\Delta^{n-1})$,

$$\sigma(z) \coloneqq \frac{1}{\sum_{j=1}^{n} \exp(\lambda z_j)} \begin{bmatrix}\exp(\lambda z_1) \\ \vdots \\ \exp(\lambda z_n)\end{bmatrix}, \lambda > 0 \tag{1}$$

where λ is referred to as the *inverse temperature constant*[1]. When $\lambda = 1$, we refer to (1) as the *standard softmax function*.

As $\lambda \to 0$, the output of the softmax function converges pointwise to the center of the simplex, i.e. a uniform probability distribution, and as $\lambda \to \infty$, the output of the softmax function converges pointwise to $e_j \in \mathbb{R}^n$, where $j = \underset{1 \leq i \leq n}{\operatorname{argmax}} e_i^T z$, i.e. all probability is placed on the largest entry of z [17], [31].

More commonly, the softmax function is referenced in terms of the individual components of (1):

$$\sigma_i(z) \coloneqq \frac{\exp(\lambda z_i)}{\sum_{j=1}^{n} \exp(\lambda z_j)} \tag{2}$$

In $\mathbb{R}^2$, (2) reduces to the *logistic* function[2] in terms of $z_i - z_j$,

$$\sigma_i(z) = \frac{\exp(\lambda z_i)}{\exp(\lambda z_i) + \exp(\lambda z_j)} = \frac{1}{1 + \exp(-\lambda(z_i - z_j))}, j \neq i \tag{3}$$

We note that (2) can be equivalently given by:

$$\sigma_i(z) = \exp\left(\lambda z_i - \log\left(\sum_{j=1}^{n} \exp(\lambda z_j)\right)\right) \tag{4}$$

While (4) is seldom used as a representation of the softmax function, in [47], it was noted that (4) represents an *exponential family*, which is the solution of the *replicator dynamics* in evolutionary game theory [2], [13], [14].

Another important representation of the softmax function is the regularized "argmax function" of a linear function under entropy regularization[3]. Let $z \in \mathbb{R}^n$, and consider the argmax of $x^T z$ given by:

$$Q(z) \coloneqq \underset{x \in \Delta^{n-1}}{\operatorname{argmax}} x^T z \tag{5}$$

[1]Depending on the context, λ is also referred to as the gain parameter [29], rationality level [16], noise level [2, pp. 128], [17], etc.

[2]Recall that a logistic function is given by $z \mapsto \frac{1}{1+e^{-\lambda z}}, \lambda, z \in \mathbb{R}$

[3]As pointed out in [4, pp. 182], the softmax function is a soft approximation of the argmax function $z \mapsto \underset{x \in \Delta^{n-1}}{\operatorname{argmax}} x^T z$, not of the max function.

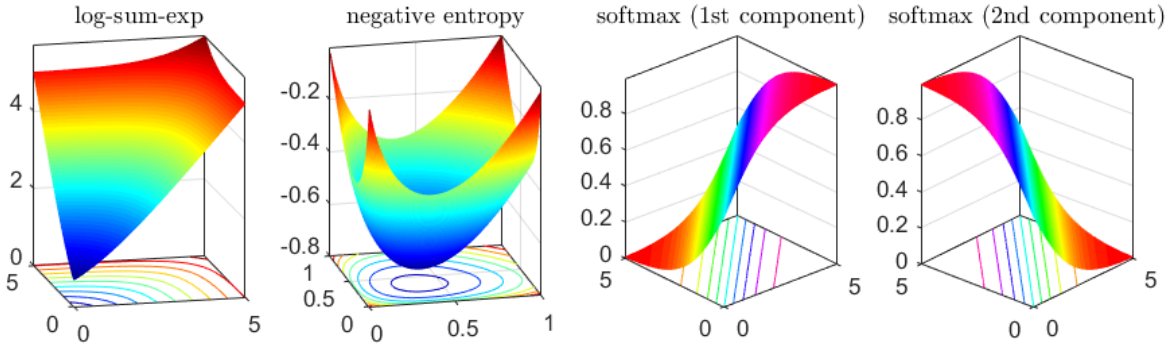

Fig. 2: Plots of the log-sum-exp, negative entropy and both components of softmax function over $\mathbb{R}^2$ with $\lambda = 1$

Suppose that $z^* \in \mathbb{R}^n$ is such that $z_j^* > z_i^*, 1 \leq i \leq n, i \neq j$, then it is clear that $x = Q(z^*) = e_j$, where $j = \underset{1 \leq i \leq n}{\operatorname{argmax}} e_i^T z^*$. In other words, (5) is maximized at the vertices of the simplex. However, (5) is set-valued. To see this, simply consider the case when two or more components of z are equal.

For many applications, it is highly desirable for $Q(z)$ to be singled-valued [16], [17], [35], [43], [44]. The most common approach to achieve this is by employing a so-called *regularizer* function[4] ψ to (5), which yields the *regularized argmax function*:

$$Q(z) = \underset{x \in \Delta^{n-1}}{\operatorname{argmax}} \left[x^T z - \psi(x) \right] \tag{6}$$

A common choice of the regularizer is the the *negative entropy*, which under the convention $0 \log(0) = 0$, is given by $\psi : \mathbb{R}^n \to \mathbb{R} \cup \{+\infty\}$,

$$\psi(x) \coloneqq \lambda^{-1} \sum_{j=1}^{n} x_j \log(x_j), \lambda > 0, x \in \Delta^{n-1} \tag{7}$$

and $+\infty$ everywhere else. When $\lambda = 1$, we refer to (7) as the *standard negative entropy* function.

Since negative entropy is λ^{-1}-strongly convex[5] in $\|\cdot\|_1$ over $\operatorname{int}(\Delta^{n-1})$ [42], by strong concavity of the argument of (6), it can be shown either by taking the gradient of the argument of (6) or by invoking the Karush-Kuhn-Tucker (KKT) conditions [6], that the unique maximizer of (6) is the softmax function evaluated at z, i.e.

$$\underset{x \in \Delta^{n-1}}{\operatorname{argmax}} \left[x^T z - \lambda^{-1} \sum_{j=1}^{n} x_j \log(x_j) \right] = \sigma(z) \tag{8}$$

It has been noted in [33], [45], [46], [50], that the objective function of (8):

$$x^T z - \lambda^{-1} \sum_{j=1}^{n} x_j \log(x_j) \tag{9}$$

represents the so called "free energy" in statistical thermodynamics. In light of this connection, in game theory terms the softmax function can be thought of as providing the *mixed strategy* with the maximum entropy which maximizes the payoff of a game [50]. We note that (8) represents an important step in Follow-the-Regularized-Leader algorithm for online linear optimization [43, pp. 127], [44, pp. 73]. Finally, the right hand side of (8) is the *Legendre-Fenchel transform* of the negative entropy function [18, pp. 102], given by $\operatorname{lse} : \mathbb{R}^n \to \mathbb{R} \cup \{+\infty\}$,

$$\operatorname{lse}(z) \coloneqq \lambda^{-1} \log \left(\sum_{j=1}^{n} \exp(\lambda z_j) \right), \lambda > 0 \tag{10}$$

When $\lambda = 1$, we refer to (10) as the *standard log-sum-exp* function. It is well known that the log-sum-exp is an approximation to the *vector-max* function $\operatorname{vecmax}(z) \coloneqq \max\{z_1, \ldots, z_n\}$ [6, pp. 72], [18, pp. 27], that is, for any $z \in \mathbb{R}^n$,

$$\operatorname{vecmax}(z) \leq \operatorname{lse}(z) \leq \operatorname{vecmax}(z) + \lambda^{-1} \log(n) \tag{11}$$

where (11) can be shown by using:

$$\exp(\lambda \operatorname{vecmax}(z)) \leq \sum_{j=1}^{n} \exp(\lambda z_j) \leq n \exp(\lambda \operatorname{vecmax}(z)) \tag{12}$$

It is for this reason that sometimes the log-sum-exp is referred to as the "softmax" function in the literature. We note that the dual or convex conjugate of the log-sum-exp function is the negative entropy [18, pp. 482], [46]. Thus $x \in \Delta^{n-1}$ can be thought of as the *dual variable* of $z \in \mathbb{R}^n$ and vice versa related through the *dual pairing* $\langle \cdot, \cdot \rangle : \Delta^{n-1} \subset \mathbb{R}^n \times \mathbb{R}^n \to \mathbb{R}$ [17]. We illustrate the log-sum-exp function as well as the negative entropy and the softmax function in Figure 2. By *Fenchel-Young* inequality, the log-sum-exp function is bounded

[4]Depending on the context, the regularizer is also referred to as admissible deterministic perturbation [2, pp. 189], penalty function [16], [17], smoothing function [38] or Bregman function [42]. For detailed construction of the regularizer, see [16], [17], [41].

[5]Recall that a function f is μ-strongly convex if there exists $\mu > 0$, such that $f(\theta z + (1-\theta)z') \leq \theta f(z) + (1-\theta) f(z') - \frac{\mu}{2} \theta (1-\theta) \|z - z'\|_2^2$ for all $z, z' \in \operatorname{dom} f$ and $\theta \in [0, 1]$. f is μ-strongly concave if $-f$ is μ-strongly convex.

below by a linear function:

$$\text{lse}(z) \geq x^T z - \psi(x), \forall x \in \Delta^{n-1}, z \in \mathbb{R}^n \tag{13}$$

Further consequences of the duality between the negative entropy and the log-sum-exp function as well as its role in game theory will not be explored at this time. Interested readers may refer to [32], [46] or any standard textbooks on convex analysis, for example, [6], [18], [22].

Finally, let $\epsilon_i, i \in \{1, \ldots, n\}$ be independent and identically distributed (i.i.d.) random variables with a *Gumbel distribution* with the cumulative distribution function:

$$\mathbb{P}(\xi_i \leq c) = \exp(-\exp(-\lambda c - \gamma)) \tag{14}$$

where $\gamma \approx 0.57721$ is the Euler-Mascheroni constant. It can be shown that [2, pp. 194],

$$\mathbb{P}\left(i = \underset{1 \leq j \leq n}{\text{argmax}}\, z_j + \epsilon_j\right) = \sigma_i(z) \tag{15}$$

In game theory terms, (15) represents the probability of choosing the *pure strategy* that maximizes the payoff or score $z \in \mathbb{R}^n$, after the payoff or score has been perturbed by a random stochastic perturbation.

A. Basic Properties of the Softmax

We briefly note some basic properties of the softmax function. First, σ maps the origin of $\mathbb{R}^n$ to be barycenter of Δ^{n-1}, that is, $\sigma(0) = n^{-1}\mathbf{1}$. We know that σ is not injective, as it can easily be shown that for any $z, z + c\mathbf{1} \in \mathbb{R}^n$, where $c \in \mathbb{R}$, we have that $\sigma(z + c\mathbf{1}) = \sigma(z)$. By definition, $\|\sigma(z)\|_1 = \sigma(z)^T\mathbf{1} = 1, \forall z \in \mathbb{R}^n$.

In a recent paper [37], it was noted that $\sigma(\mathbf{P}(z)) = \mathbf{P}\sigma(z)$, where $\mathbf{P}$ is any permutation matrix, and that the standard softmax function satisfies a type of "coordinate non-expansiveness" property, whereby if $z_j \geq z_i$, then $0 \leq \sigma_j(z) - \sigma_i(z) \leq \frac{1}{2}(z_j - z_i)$. The last property can be derived by exploiting the properties of the hyperbolic tangent function. It was also noted that these properties of the softmax function bears similarities with the Euclidean projection onto Δ^{n-1} [37].

In addition, in a direction that is tangential to the aim of this paper, [34] is interested in finding a bound on the softmax function. It can be shown that,

$$\sigma_i(z) = \frac{\exp(\lambda z_i)}{\sum_{j=1}^{n} \exp(\lambda z_j)} \geq \prod_{\substack{j=1 \\ j \neq i}}^{n} \frac{1}{1 + \exp(-\lambda(z_i - z_j))} \tag{16}$$

(16) is referred in "one-vs-each" bound which can be generalized to bounds on arbitrary probabilities [34]. We note that this inequality is tight for $n = 2$, see (3).

IV. Review of Convex Optimization and Monotone Operator Theory

In this section we review some of the definitions and results from convex optimization and monotone operator theory. Most of the proofs to the propositions in this section can be found in standard references such as [6], [18], [19], [20], [21], [22]. Throughout this section we will assume that $\mathbb{R}^n$ is equipped with the standard inner product $\langle z, z' \rangle := \sum_{i=1}^{n} z_i z_i'$ with the induced 2-norm $\|z\|_2 := \sqrt{\langle z, z \rangle}$. We assume the domain of f, $\text{dom}\, f$, is convex.

Definition 2. A function $f : \text{dom}\, f \subseteq \mathbb{R}^n \to \mathbb{R}$ is convex if

$$f(\theta z + (1-\theta)z') \leq \theta f(z) + (1-\theta) f(z') \tag{17}$$

for all $z, z' \in \text{dom}\, f$ and $\theta \in [0, 1]$ and strictly convex if (17) holds strictly whenever $z \neq z'$ and $\theta \in (0, 1)$

The convexity of a C^2 function f is easily determined through its Hessian $\nabla^2 f$.

Proposition 1. *Let f be C^2. Then f is convex if and only if $\text{dom}\, f$ is convex and its Hessian is positive semidefinite, that is, for all $z, v \in \text{dom}\, f$,*

$$v^T \nabla^2 f(z) v \geq 0 \tag{18}$$

and strictly convex if $\nabla^2 f(z)$ is positive definite for all $z \in \text{dom}\, f$

Next, we will introduce the concept of a monotone operator and related properties. A monotone operator is usually taken as a set-valued relation, however, it is also natural for the definitions related to a monotone operator to be directly applied to single-valued maps. A widely used and standard reference for the following definitions is [20].

Definition 3. ([20, pp. 154]) An operator (or mapping) $F : \mathcal{D} \subseteq \mathbb{R}^n \to \mathbb{R}^n$ is said to be

- pseudo monotone on $\mathcal{D}$ if

$$F(z')^T(z - z') \geq 0 \implies F(z)^T(z - z') \geq 0, \forall z, z' \in \mathcal{D} \tag{19}$$

- pseudo monotone plus on $\mathcal{D}$ if it is pseudo monotone on $\mathcal{D}$ and

$$\begin{aligned} &F(z')^T(z - z') \geq 0 \text{ and } F(z)^T(z - z') = 0 \\ \implies &F(z) = F(z'), \forall z, z' \in \mathcal{D} \end{aligned} \tag{20}$$

- monotone on $\mathcal{D}$ if

$$(F(z) - F(z'))^T(z - z') \geq 0, \forall z, z' \in \mathcal{D} \tag{21}$$

- monotone plus on $\mathcal{D}$ if it is monotone on $\mathcal{D}$ and

$$(F(z) - F(z'))^T(z - z') = 0 \implies F(z) = F(z'), \forall z, z' \in \mathcal{D} \tag{22}$$

- strictly monotone on $\mathcal{D}$ if

$$(F(z) - F(z'))^T(z - z') > 0, \forall z, z' \in \mathcal{D}, z \neq z' \tag{23}$$

Clearly, a strictly monotonicity implies monotone plus, which in turn implies monotone, pseudo monotone plus and pseudo monotone. By definition, every strictly monotone operator is an injection. We will refer to an operator F as being *(strictly) anti-monotone* if $-F$ is (strictly) monotone.

The following proposition provides a natural connection between C^1, convex functions and monotone gradient maps.

Proposition 2. *A C^1 function f is convex if and only if*

$$(\nabla f(z) - \nabla f(z'))^T (z - z') \geq 0, \forall z, z' \in \operatorname{dom} f \tag{24}$$

and strictly convex if and only if

$$(\nabla f(z) - \nabla f(z'))^T (z - z') > 0, \forall z, z' \in \operatorname{dom} f, z \neq z' \tag{25}$$

Next, we will introduce the notions of Lipschitz continuity and co-coercivity and show that the two concepts can be related through the gradient of a convex function.

Definition 4. An operator (or mapping) $F : \mathcal{D} \subseteq \mathbb{R}^n \to \mathbb{R}^n$ is said to be

- *Lipschitz* (or *L-Lipschitz*) if there exists a $L > 0$ such that

$$\|F(z) - F(z')\|_2 \leq L\|z - z'\|_2, \forall z, z' \in \mathcal{D} \tag{26}$$

  If $L = 1$ in (26), then F is referred to as *nonexpansive*. Otherwise, if $L \in (0, 1)$, then F is referred to as *contractive*.
- *co-coercive* (or $\frac{1}{L}$*-co-coercive*) if there exists a $L > 0$ such that

$$(F(z) - F(z')^T (z - z') \geq \frac{1}{L}\|F(z) - F(z)'\|_2^2, \forall z, z' \in \mathcal{D} \tag{27}$$

  If $L = 1$ in (27), then F is referred to as *firmly nonexpansive*.

By Cauchy-Schwarz inequality, every $\frac{1}{L}$-co-coercive operator is L-Lipschitz, in particular, every firmly nonexpansive operator is nonexpansive. However, the reverse need not be true, for example $f(z) = -z$ is nonexpansive but not firmly nonexpansive. Fortunately, a theorem exists which provides the condition for when a L-Lipschitz operator is also $\frac{1}{L}$-co-coercive. This is known as the Baillon-Haddad theorem in convex analysis [23].

Theorem 1. *(Baillon-Haddad theorem)*

Let $f : \operatorname{dom} f \subseteq \mathbb{R}^n \to \mathbb{R}$ be a C^1, convex function on $\operatorname{dom} f$ and such that ∇f is L-Lipschitz continuous for some $L > 0$, then ∇f is $\frac{1}{L}$-co-coercive.

Proof. The proof of Baillon-Haddad theorem can be found in [21, pp. 40], Theorem 3.13. ■

Finally, we will introduce the notion of *maximal monotonicity*. Let $H : \mathbb{R}^n \to 2^{\mathbb{R}^n}$ be the set-valued map, where $2^{\mathbb{R}^n}$ is the *power set* of $\mathbb{R}^n$. Let graph of H be given by $\text{gra}H := \{(u, v) \in \mathbb{R}^n \times \mathbb{R}^n | v = Hu\}$. H is said to be *monotone* if $(u - u')^T (v - v') \geq 0, v \in H(u), v' \in H(u')$.

Definition 5. ([19, pp. 297]) Let $H : \mathbb{R}^n \to 2^{\mathbb{R}^n}$ be monotone. Then H is maximal monotone if there exists no monotone operator $G : \mathbb{R}^n \to 2^{\mathbb{R}^n}$ such that $\operatorname{gra} G$ properly contains $\operatorname{gra} H$, i.e., for every $(u, v) \in \mathbb{R}^n \times \mathbb{R}^n$,

$$(u, v) \in \operatorname{gra} H \Leftrightarrow (\forall (u', v') \in \operatorname{gra} H)\,(u - u')^T (v - v') \geq 0 \tag{28}$$

By Zorn's Lemma, every monotone operator can be extended to a maximal monotone operator [18, pp. 535], [19, pp. 297]. For the scope of this paper, we are interested when a single-valued map is maximal monotone. The following proposition provides a simple characterization of this result [18, pp. 535].

Proposition 3. *If a continuous mapping $F : \mathbb{R}^n \to \mathbb{R}^n$ is monotone, it is maximal monotone. In particular, every differentiable monotone mapping is maximal monotone.*

V. Derivation of Properties of Softmax function

In this section we will derive several properties of the softmax function using tools from convex analysis and monotone operator theory. We will start by establishing the connection between the log-sum-exp function and the softmax function.

It has long been known that the softmax function is the gradient map of a convex potential function [31], however, the fact that its potential function is the log-sum-exp function is rarely discussed[6]. We make this connection clear with the following proposition.

Proposition 4. *The softmax function is the gradient of the log-sum-exp function, that is, $\sigma(z) = \nabla \operatorname{lse}(z)$.*

Proof. Evaluating the partial derivative of lse at each component of $z_i, i \in \{1, \ldots, n\}$ yields $\dfrac{\partial \operatorname{lse}(z)}{\partial z_i} = \dfrac{\exp(\lambda z_i)}{\sum_{j=1}^n \exp(\lambda z_j)}$. By definition of the gradient of lse, we have:

$$\nabla \operatorname{lse}(z) = \begin{bmatrix} \dfrac{\partial \operatorname{lse}(z)}{\partial z_1} \\ \vdots \\ \dfrac{\partial \operatorname{lse}(z)}{\partial z_n} \end{bmatrix} = \frac{1}{\sum\limits_{j=1}^n \exp(\lambda z_j)} \begin{bmatrix} \exp(\lambda z_1) \\ \vdots \\ \exp(\lambda z_n) \end{bmatrix} = \sigma(z)$$

■

Next, we calculate the Hessian of the log-sum-exp function (and hence the Jacobian of softmax).

Proposition 5. *The Jacobian of the softmax function and Hessian of the log-sum-exp function is given by:*

$$\mathbf{J}[\sigma(z)] = \nabla^2 \operatorname{lse}(z) = \lambda(\operatorname{diag}(\sigma(z)) - \sigma(z)\sigma(z)^T) \tag{29}$$

where (29) *is a symmetric positive semidefinite matrix and satisfies $\mathbf{J}[\sigma(z)]\mathbf{1} = \mathbf{0}$, that is, $\mathbf{1}$ is the eigenvector associated with the zero eigenvalue of $\mathbf{J}[\sigma(z)]$.*

Proof. The diagonal entries of $\nabla^2 \operatorname{lse}$ are given by,

$$\frac{\partial^2 \operatorname{lse}(z)}{\partial z_i^2} = \frac{\lambda \left[\exp(\lambda z_i) \sum\limits_{j=1}^n \exp(\lambda z_j) - \exp(\lambda z_i)^2 \right]}{\left(\sum\limits_{j=1}^n \exp(\lambda z_j) \right)^2}$$

[6] Although not explicitly stated, this relationship could also be found in [6, pp. 93] and various other sources.

and the off-diagonal entries of $\nabla^2 \operatorname{lse}$ are given by the mixed partials,

$$\frac{\partial^2 \operatorname{lse}(z)}{\partial z_k \partial z_i} = \frac{-\lambda \exp(\lambda z_k)\exp(\lambda z_i)}{\Big(\sum\limits_{j=1}^{n} \exp(\lambda z_j)\Big)^2}$$

Assembling the partial derivatives, we obtain the Hessian of lse and the Jacobian of σ:

$$\mathbf{J}[\sigma(z)] = \nabla^2 \operatorname{lse}(z) = \lambda(\operatorname{diag}(\sigma(z)) - \sigma(z)\sigma(z)^T) \quad (30)$$

The symmetry of $\mathbf{J}[\sigma(z)]$ comes from the symmetric structure of the diagonal and outer product terms. The positive semi-definiteness of $\mathbf{J}[\sigma(z)]$ follows from an application of the Cauchy-Schwarz inequality [6, pp. 74]. It can be shown through direct computation that $\mathbf{J}[\sigma(z)]\mathbf{1} = \mathbf{0}$, or refer to [2, pp. 213]. ■

This result is known early on in [31], [32] and can be found in [2, pp. 195][6, pp. 74]. As a trivial consequence of Proposition 5, we can write the individual components of $\mathbf{J}[\sigma(z)]$ as:

$$\mathbf{J}_{ij}[\sigma(z)] = \lambda\sigma_i(z)(\delta_{ij} - \sigma_j(z)) \quad (31)$$

where $\delta_{ij} = 1$, if $i = j$, and $\delta_{ij} = 0$, if $i \neq j$ is the Kronecker delta. This representation is preferred for machine learning related applications.

Finally, let $x = \sigma(z)$, then we have:

$$\nabla^2 \operatorname{lse}(z)\big|_{x=\sigma(z)} = \lambda(\operatorname{diag}(x) - xx^T) \quad (32)$$

We note this is precisely the matrix term appearing in the replicator dynamics [2, pp. 229], [39], that is,

$$\dot{x} = \nabla^2 \operatorname{lse}(z)\big|_{x=\sigma(z)} u = \lambda(\operatorname{diag}(x) - xx^T)u \quad (33)$$

where $u \in \mathbb{R}^n$ is a *payoff vector*. This result provides an important connection between the field of evolutionary game theory with convex analysis and monotone operator theory.

Proposition 6. *The log-sum-exp function is C^2, convex and not strictly convex on $\mathbb{R}^n$.*

Proof. The convexity of the log-sum-exp function follows from Proposition 5. This result can also be established by using the AM-GM inequality [24, pp. 14], or simply by recognizing that the log-sum-exp function the convex conjugate of the negative entropy [18, pp. 102].

To show that log-sum-exp is not strictly convex, let $z = z_o + c\mathbf{1}, z_o \in \mathbb{R}^n, c \in \mathbb{R}$, then,

$$\operatorname{lse}(z) = \operatorname{lse}(z_o + c\mathbf{1}) = \operatorname{lse}(z_o) + c$$

We have that lse is affine along the line given by $z = z_o + c\mathbf{1}$, thus the log-sum-exp function is not strictly convex. This result is also noted in [18, pp. 48]. ■

Corollary 1. *The softmax function is monotone, that is,*

$$(\sigma(z) - \sigma(z'))^T(z - z') \geq 0, \forall z, z' \in \mathbb{R}^n \quad (34)$$

and not strictly monotone on $\mathbb{R}^n$.

Proof. Monotonicity of σ follows directly from the convexity of the log-sum-exp function. Since the log-sum-exp function is not strictly convex on $\mathbb{R}^n$, therefore by Proposition 2, σ fails to be strictly monotone. Alternatively, since every strictly monotone operator is injective, therefore σ is not strictly monotone on $\mathbb{R}^n$. ■

The monotonicity of σ allows us to state a stronger result.

Corollary 2. *The softmax function is a maximal monotone operator, that is, there exists no monotone operator such that its graph properly contains the graph of the softmax function.*

Proof. This directly follows from σ being a continuous, monotone map, see Proposition 3. ■

Next, we show that under appropriate conditions, the softmax function is a contraction in $\|\cdot\|_2$.

Lemma 1. *(Lemma 55 [51]) A C^2, convex function $f : \mathbb{R}^n \to \mathbb{R}$ has a Lipschitz continuous gradient with Lipschitz constant $L > 0$ if for all $z, v \in \mathbb{R}^n$*

$$0 \leq v^T \nabla^2 f(z) v \leq L\|v\|^2 \quad (35)$$

The proof of the above lemma can be derived through a non-trivial application of Fenchel duality or as a consequence of the Baillon-Haddad theorem, see Theorem 1.

Proposition 7. *The softmax function is L-Lipschitz with respect to $\|\cdot\|_2$ with $L = \lambda$, that is, for all $z, z' \in \mathbb{R}^n$,*

$$\|\sigma(z) - \sigma(z')\|_2 \leq \lambda\|z - z'\|_2 \quad (36)$$

where λ is the inverse temperature constant.

Proof. Given the Hessian of lse in Proposition 5, we have for all $z, v \in \mathbb{R}^n$

$$v^T \nabla^2 \operatorname{lse}(z) v = \lambda \left(\sum_{i=1}^{n} v_i^2 \sigma_i(z) - \left(\sum_{i=1}^{n} v_i \sigma_i(z) \right)^2 \right) \quad (37)$$

Since the second term on the right hand side of (37) is nonnegative, therefore

$$\begin{aligned} v^T \nabla^2 \operatorname{lse}(z) v \leq & \lambda \sum_{i=1}^{n} v_i^2 \sigma_i(z) \leq \lambda \sup\{\sigma_i(z)\} \sum_{i=1}^{n} v_i^2 \\ \implies & v^T \nabla^2 \operatorname{lse}(z) v \leq \lambda \|v\|_2^2 \end{aligned} \quad (38)$$

where $\sup\{\sigma_i(z)\} = 1, \forall i \in \{1, \ldots, n\}, \forall z \in \mathbb{R}^n$. By Proposition 6, $\nabla^2 \operatorname{lse}(z)$ is positive semidefinite. Hence using (38), we have,

$$0 \leq v^T \nabla^2 \operatorname{lse}(z) v \leq \lambda \|v\|_2^2 \quad (39)$$

By Lemma 1, σ is Lipschitz with $L = \lambda$. ■

We note that Proposition 7 can also be established using Theorem 4.2.1. in [22, pp. 240], which resorts to using duality between the negative entropy and the log-sum-exp function.

As a minor consequence of Proposition 7, by Cauchy-Schwarz inequality, we have:

$$(\sigma(z) - \sigma(z'))^T(z - z') \leq \lambda \|z - z'\|_2^2 \quad (40)$$

Corollary 3. *The softmax function is $\frac{1}{L}$-co-coercive with respect to $\|\cdot\|_2$ with $L = \lambda$, that is, for all $z, z' \in \mathbb{R}^n$,*

$$(\sigma(z) - \sigma(z'))^T (z - z') \geq \frac{1}{\lambda} \|\sigma(z) - \sigma(z')\|_2^2 \tag{41}$$

where λ is the inverse temperature constant.

Proof. Follows directly from Baillon - Haddad Theorem, see Theorem 1. ∎

Proposition 7 and Corollary 3 show that the inverse temperature constant λ is crucial in determining the Lipschitz and co-coercive properties of the softmax function. We summarize these properties with respect to $\|\cdot\|_2$ in the following corollary.

Corollary 4. *The softmax function is:*

- *λ-Lipschitz and $\frac{1}{\lambda}$-co-coercive for $\lambda > 1$*
- *Nonexpansive and firmly nonexpansive for $\lambda = 1$*
- *Contractive for $\lambda \in (0, 1)$*

where λ is the inverse temperature constant.

Finally, we will state an additional consequence of σ being Lipschitz, monotone operator with a symmetric Jacobian matrix over all $\mathbb{R}^n$.

Corollary 5. *The softmax function is monotone plus, pseudo monotone plus and pseudo monotone on $\mathbb{R}^n$.*

Proof. This follows from the chain of implications in [20, pp. 164]. ∎

VI. Application to Population Games and Reinforcement Learning

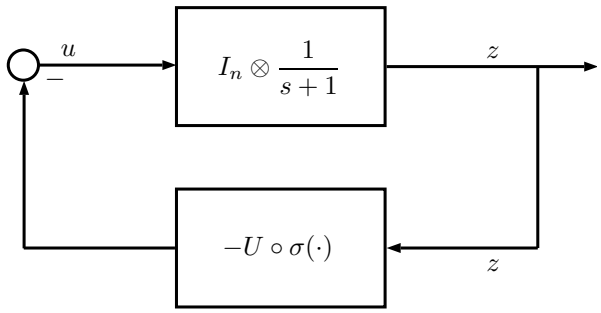

Fig. 3: Feedback representation of the exponentially-discounted reinforcement learning scheme under the softmax selection rule

In this section we demonstrate the application of these new properties in the context of population games and stateless continuous-time score-based reinforcement learning. For related work in this direction, see [16], [17], [40]. For background on population games, see [2]. Consider a population game $\mathcal{G}$ with a single population[7], strategy set $\mathcal{A} = \{1, \ldots, n\}$ and payoff function $U : \mathbb{R}^n \to \mathbb{R}^n$. We assume that the population aggregates the raw payoff $u \in \mathbb{R}^n$ from a game $\mathcal{G}$ into score variables $z \in \mathbb{R}^n$ via the rule:

$$z_i = \int_0^t e^{-(t-\tau)} u_i(\tau) \mathrm{d}\tau \tag{42}$$

where $u_i = U_i(x) \in \mathbb{R}$ is the payoff to the i^{th} strategy, $x \in \Delta^{n-1}$ is the mixed strategy, and $z_i \in \mathbb{R}$ is the score variable associated with the i^{th} strategy. (42) is known as *exponentially-discounted learning scheme*, under which the population allocates exponentially more weight to recent observations of the payoff [16], [17]. Taking the time derivative of (42) yields the *score dynamics*:

$$\dot{z}_i = u_i - z_i \tag{43}$$

We refer to (43) as the *exponentially-discounted score dynamics*. This form of score dynamics was investigated in [16], [48], [50], [52]. We can express (43) in vector form as:

$$\dot{z} = u - z \tag{44}$$

Suppose that the score variable z is mapped to the strategy x through the softmax selection rule, i.e. $x = \sigma(z)$, then the payoff can be written as $u = U(x) = U(\sigma(z))$. Expressing the composition between the softmax selection rule with the payoff function as $U \circ \sigma(z) \coloneqq U(\sigma(z))$, then we can write (44) as

$$\dot{z} = (U \circ \sigma)(z) - z \tag{45}$$

The overall learning scheme can be represented as a closed-loop feedback system in Figure 3, where I_n is the $n \times n$ identity matrix, $\otimes$ is the Kronecker product, and $\frac{1}{s+1}$ is the transfer function of (43) from u_i to z_i. Under control-theoretic convention, we represent the feedback subsystem using negative of the composite map i.e. $-U \circ \sigma$. The closed-loop system is equivalently represented by:

$$\begin{cases} \dot{z} = u - z \\ u = U(x) \\ x = \sigma(z) \end{cases} \tag{46}$$

Next, we find conditions on U and $U \circ \sigma$ under which the *induced strategy trajectory* $x(t) = \sigma(z(t))$ converges towards the *quantal response equilibrium* (or more precisely, a *logit response equilibrium*) of the game [11][16]. For backgrounds on dynamical system and Lyapunov theory, see [49]. This analysis is largely inspired by [52] but performed in a population games setup.

We observe that the equilibria of the overall closed-loop system (46) are the fixed points of the map $z \mapsto (U \circ \sigma)(z)$. This fixed-point condition can be restated as:

$$\begin{cases} \dot{z} = 0 \implies u^* = z^* \\ u^* = U(x^*) \\ x^* = \sigma(z^*) \end{cases} \tag{47}$$

From (47), it is clear that the overall score map σ is

[7] We can identify such the action of a population with the action of a player in a "mean-field" sense, see related [16], [17], [40].

invertible over the set of rest points of (45) with the inverse U.

Proposition 8. *(Proposition 2 [52]) Let $\mathcal{E}$ denote the set of rest points for* (45)*, i.e.*

$$\mathcal{E} = \{z^* \in \mathbb{R}^n | z^* = (U \circ \sigma)(z^*)\} \tag{48}$$

The map σ is one-to-one over the set $\mathcal{E}$.

We note that the existence of the fixed point is guaranteed by the Brouwer's Fixed Point Theorem provided that $U \circ \sigma$ is a continuous function with bounded range [52]. Since $z^* = u^*$, therefore the fixed point z^* is mapped through σ to a quantal response equilibrium.

Proposition 9. *$x^* = \sigma(z^*) = \sigma(u^*)$ is the quantal response equilibrium of the game $\mathcal{G}$.*

In the following, we show induced strategy trajectory $x(t) = \sigma(z(t))$ converges towards quantal response equilibrium x^* under different assumptions on U or $U \circ \sigma$.

Assumption 1. $U \circ \sigma$ is $\|\cdot\|_\infty$-contractive, that is, there exists a constant $L \in (0, 1)$ such that for all score variables $z, z' \in \mathbb{R}^n$,

$$\|(U \circ \sigma)(z) - (U \circ \sigma)(z')\|_\infty \leq L\|z - z'\|_\infty \tag{49}$$

Proposition 10. *(Theorem 4, [52]) Under Assumption 1, the unique fixed point $z^* \in \mathbb{R}^n$ of the composite map $U \circ \sigma : \mathbb{R}^n \to \mathbb{R}^n$ is globally asymptotically stable for* (45)*.*

Example 1. By equivalence of norms [6, pp. 636], σ is $\|\cdot\|_\infty$-contractive if $\sqrt{n}\lambda < 1$. Then for any game where payoff function U is a $\|\cdot\|_\infty$-contraction, then $U \circ \sigma$ is a $\|\cdot\|_\infty$-contraction. Proposition 10 implies that the induced strategy trajectory $x(t) = \sigma(z(t))$ converges to the quantal response equilibrium x^*.

Assumption 2. $U \circ \sigma$ is anti-monotone, that is, for all $z, z' \in \mathbb{R}^n$,

$$(z - z')^T((U \circ \sigma)(z) - (U \circ \sigma)(z')) \leq 0 \tag{50}$$

Theorem 2. *Under Assumption 2, any fixed point $z^* \in \mathbb{R}^n$ of the composite map $U \circ \sigma : \mathbb{R}^n \to \mathbb{R}^n$ is globally asymptotically stable for* (45)*.*

Proof. Consider the Lyapunov function:

$$V(z) = \frac{1}{2}\|z - z^*\|_2^2 \tag{51}$$

Taking the time derivative of $V(z)$ along the solution of (45) yields:

$$\begin{aligned}\dot{V}(z) &= \nabla V(z)^T \dot{z} = (z - z^*)^T((U \circ \sigma)(z) - z)) \\ &= (z - z^*)^T((U \circ \sigma)(z) - (U \circ \sigma)(z^*)) - \|z - z^*\|_2^2\end{aligned}$$

where the last line was obtained by adding and subtracting z^* and using the equilibrium condition in (47). By (50), $\dot{V}(z) \leq 0, \forall z \in \mathbb{R}^n$ and $\dot{V}(z) = 0$ only if $(z - z^*)^T((U \circ \sigma)(z) - (U \circ \sigma)(z^*)) = 0$ and $\|z - z^*\|_2^2 = 0$. Since the latter is true only if $z = z^*$, therefore by LaSalle's theorem [49, pp. 128], z^* is globally asymptotically stable for (45). ∎

Example 2. Consider matching in pure coordination games [2], since the payoff function is given by $U(x) = -\alpha I_n x$ where I_n is the $n \times n$ identity matrix. From the monotonicity of σ, it is easy to see that $U \circ \sigma$ is anti-monotone for all $\alpha > 0$. Theorem 2 implies induced strategy trajectory $x(t) = \sigma(z(t))$ converges to the quantal response equilibrium x^*.

Assumption 3. The payoff U is anti-monotone, that is, for all $x, x' \in \Delta^{n-1}$,

$$(x - x')^T(U(x) - U(x')) \leq 0 \tag{52}$$

Theorem 3. *Under Assumption 3, any solution of* (45) *converges to the subspace $\mathcal{M} = \{z \in \mathbb{R}^n | z = z^* + \mathbf{1}c, c \in \mathbb{R}\}$ where $z^* \in \mathbb{R}^n$ is the fixed point of $U \circ \sigma$, $z^* = U \circ \sigma(z^*)$. Moreover, $x(t) = \sigma(z(t))$ converges to the quantal response equilibrium $x^* = \sigma(z^*)$ of game $\mathcal{G}$.*

Proof. First, we show that any point in $\mathcal{M} = \{z \in \mathbb{R}^n | z = z^* + \mathbf{1}c, c \in \mathbb{R}\}$ where $z^* \in \mathbb{R}^n$ is the fixed point of $U \circ \sigma$, is a rest point of (45).

At rest, $\dot{z} = 0 \implies U \circ \sigma(z^*) = U(\sigma(z^*)) = z^*$. By non-injectivity of σ, we have $U(\sigma(z^*)) = U(\sigma(z^* + \mathbf{1}c)), \forall c \in \mathbb{R}$. Since the latter condition is satisfied by the set of points given by $\mathcal{M}$, therefore every point in $\mathcal{M}$ is a rest point of (45).

Next, consider the Lyapunov function given by a *Bregman-type*[8] function:

$$V(z) = \operatorname{lse}(z) - \operatorname{lse}(z^*) - \nabla \operatorname{lse}(z^*)^T(z - z^*) \tag{53}$$

where $z^* = U \circ \sigma(z^*)$ and lse is given in (10). Recall that by Proposition 6, lse is convex and by Proposition 4, $\nabla \operatorname{lse} = \sigma$. By convexity of lse, $V(z) \geq 0, \forall z \in \mathbb{R}^n$. Using $\|\sigma(z)\|_1 = \sigma(z)^T\mathbf{1} = 1$ and $\operatorname{lse}(z + \mathbf{1}c) = \operatorname{lse}(z) + c$, it is trivial to show that $V(z) = 0, \forall z \in \mathcal{M}$.

Taking the time derivative of $V(z)$ along the solution of (45) yields:

$$\begin{aligned}\dot{V}(z) =& \nabla V(z)^T \dot{z} = (\sigma(z) - \sigma(z^*))^T(-z + u) \\ =& (\sigma(z) - \sigma(z^*))^T(-z + z^* - z^* + u) \\ =& -(\sigma(z) - \sigma(z^*))^T(z - z^*) + (\sigma(z) - \sigma(z^*))^T(u - u^*)\end{aligned}$$

By Corollary 3, σ is co-coercive, therefore:

$$\dot{V}(z) \leq -\frac{1}{\lambda}\|\sigma(z) - \sigma(z^*)\|_2^2 + (\sigma(z) - \sigma(z^*))^T(u - u^*)$$

Since $u^* = U(\sigma(z))$, $x = \sigma(z)$, by (52), $\dot{V}(z) \leq 0, \forall z \in \mathbb{R}^n$. Thus $V(z(t))$ is non-increasing and bounded from below by 0, hence it converges as $t \to \infty$. In addition, (52) implies that $\dot{V}(z) \leq -\frac{1}{\lambda}\|\sigma(z) - \sigma(z^*)\|_2^2$, so that $\int_0^t \|\sigma(z) - \sigma(z^*)\|_2^2 \mathrm{d}\tau \leq \lambda[V(z(0)) - V(z(t))]$. Since $V(z(t))$ converges, therefore $\lim_{t\to\infty} \int_0^t \|\sigma(z(\tau)) - \sigma(z^*)\|_2^2 \mathrm{d}\tau$ exists and is finite. We can show that the *induced dynamics* $\dot{x}(t)$ is Lipschitz continuous (Theorem 4.4.2, [2, pp. 131]), hence

[8]Since lse is not strictly convex, (53) is not a true Bregman divergence [18, pp. 258], whereas (51) is a true Bregman divergence in that $\frac{1}{2}\|z\|_2^2$ is strictly convex.

$\dot{x}(t)$ is bounded, and therefore $x(t)$ is uniformly continuous. From Barbalat's lemma [49, pp. 323], we conclude that $\sigma(z(t))-\sigma(z^*) \to 0$ as $t \to \infty$. This implies that z converges to $\mathcal{M}$ and x converges to x^*, which by Proposition 9 is the quantal response equilibrium of game $\mathcal{G}$. ■

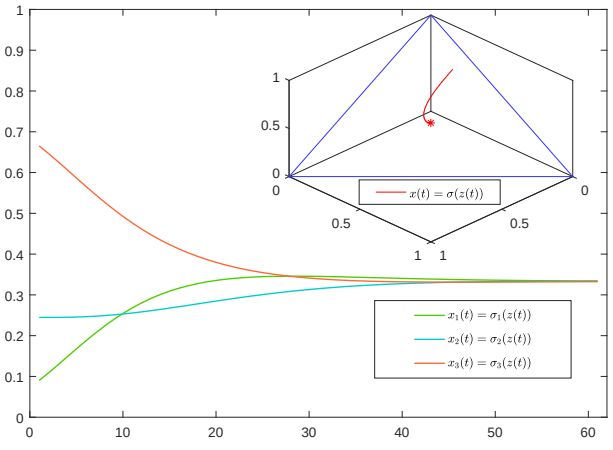

Fig. 4: Convergence of the induced strategy trajectory $x(t)$ towards the quantal response equilibrium of the standard RPS game

We note that Assumption 3 is equivalent to game $\mathcal{G}$ being a *stable game* [2, pp. 79], [15]. The representative game from the class of stable games is the standard Rock-Paper-Scissor (RPS) game given by the payoff matrix:

$$\mathbf{A} = \begin{bmatrix} 0 & -1 & 1 \\ 1 & 0 & -1 \\ -1 & 1 & 0 \end{bmatrix} \tag{54}$$

which generates the payoff function $U(x) = \mathbf{A}x$. We present a simulation of this game under the exponentially-discounted score dynamics (44) with $\lambda = 1$ and fixed step size of 0.1. The resulting induced strategy trajectory $x(t)$ is shown in Figure 4, which by Theorem 3 is guaranteed to converge to the quantal response equilibrium of the game.

VII. Conclusion and Open Problems

This paper presents a thorough analysis of the softmax function using tools from convex analysis and monotone operator theory. We have shown that the softmax function is the monotone gradient map of the log-sum-exp function and that the inverse temperature parameter λ determine the Lipschitz and co-coercivity properties of the softmax function. We have demonstrated demonstrate the usefulness of these properties through an application in game theory and reinforcement learning. We hope that researchers could adapt our results presented here and apply them to their domain-specific problems. Next, we will elaborate on some of the open problems.

While the proof of Corollary 1 presented in this paper is straightforward and unambiguous, its implications should strike the readers as counter-intuitive. To illustrate this, consider the softmax function over $\mathbb{R}^2$. Let $\phi(z) = \log(\sum_{j=1}^{n} \exp(\lambda z_j))$, then by Corollary 1 and using (4), the left-hand side of (34) can be written as:

$$\begin{aligned}&(\sigma_1(z) - \sigma_1(z'))(z_1 - z_1') + (\sigma_2(z) - \sigma_2(z'))(z_2 - z_2') \\ =&(\exp(\lambda z_1 - \phi(z)) - \exp(\lambda z_1' - \phi(z'))(z_1 - z_1') + \\ &(\exp(\lambda z_2 - \phi(z)) - \exp(\lambda z_2' - \phi(z'))(z_2 - z_2')\end{aligned}$$

Suppose for the moment being that $\phi(z), \phi(z')$ were equal to zero, then we have:

$$(\exp(\lambda z_1) - \exp(\lambda z_1'))(z_1 - z_1') + (\exp(\lambda z_2) - \exp(\lambda z_2'))(z_2 - z_2')$$

Since the exponential e^z is a monotonically increasing function, therefore each of the terms in the summation is non-negative, and the overall expression is non-negative. However, if $\phi(z), \phi(z')$ are non-zero, then it is difficult to establish the sign of individual products $(\exp(\lambda z_1 - \phi(z)) - \exp(\lambda z_1' - \phi(z')))(z_1 - z_1')$ and $(\exp(\lambda z_2 - \phi(z)) - \exp((\lambda z_2' - \phi(z')))(z_2 - z_2')$, and the sign of their summation. Despite of this ambiguity, Corollary 1 still gives that the summation of the two terms is nonnegative. The resolution of this ambiguity is equivalent to finding a direct proof of the monotonicity property of σ.

For the proof of Theorem 2, we have assumed that $U \circ \sigma$ is anti-monotone. This is equivalent to requiring the composition between σ and $-U$ to be monotone. It is not known at this point if any linear operator other than $-U(x) = \alpha I_n x, \alpha > 0$ can preserve the monotonicity property of σ. Such a linear operator would constitute as a new class of games for which the induced strategy trajectory of the score dynamics given in (43) converges to the quantal response equilibrium.

The Jacobian of the softmax function $\mathbf{J}[\sigma(z)]$ has been provided in Proposition 5 and it is known that this matrix is symmetric positive semidefinite. We know that every symmetric matrix has the eigenvalue decomposition $\mathbf{A} = \mathbf{Q}^T \mathbf{\Lambda} \mathbf{Q}$ where $\mathbf{Q}$ is an orthogonal matrix and $\mathbf{\Lambda} = \text{diag}(\lambda_1, \ldots, \lambda_n)$, where λ_i is the i^{th} eigenvalue of $\mathbf{A}$. This factorization would provide an additional characterization of the replicator dynamics.

Finally, for applications in reinforcement learning, it is desirable to use a generalized version of the softmax function given by:

$$\sigma_i(z) = \frac{\exp(\lambda_i z_i)}{\sum_{j=1}^{n} \exp(\lambda_j z_j)} \tag{55}$$

Here, each strategy i is associated with an inverse temperature constant $\lambda_i > 0$, which can be adjusted independently to improve an agent's learning performance. The relationship between the individual parameters λ_i with the convergence properties of score dynamics under the choice rule given by (55) has been investigated in [52] but is not yet fully characterized at this point.

REFERENCES

[1] H. Young and S. Zamir, Handbook of Game Theory, Volume 4, 1st ed. Amsterdam: Elsevier, North-Holland, 2015.

[2] W. H. Sandholm, Population Games and Evolutionary Dynamics. Cambridge, MA, USA: MIT Press, 2010.

[3] R. Sutton and A. Barto, Reinforcement Learning: An Introduction. Cambridge, MA, USA: MIT Press, 1998.

[4] I. Goodfellow, Y. Bengio, and A. Courville, Deep Learning. Cambridge, MA, USA: MIT Press, 2016.

[5] C. M. Bishop, Pattern Recognition and Machine Learning. Secaucus, NJ, USA: Springer, 2006.

[6] S. Boyd and L. Vandenberghe, Convex optimization, 1st ed. Cambridge, UK: Cambridge University Press, 2004.

[7] D. Bloembergen, K. Tuyls, D. Hennes, and M. Kaisers, "Evolutionary dynamics of multi-agent learning: A survey", J. Artif. Intell. Res., vol. 53, no. 1, pp. 659-697, May 2015.

[8] G. Weiss, Multiagent Systems, 2nd ed. Cambridge, MA, USA: MIT Press, 2013.

[9] R. Luce, Individual Choice Behavior: A Theoretical Analysis. NY, Wiley, 1959.

[10] J. Bridle, "Probabilistic Interpretation of Feedforward Classification Network Outputs, with Relationships to Statistical Pattern Recognition", Neurocomputing: Algorithms, Architectures and Applications, F. Soulie and J. Herault, eds., pp. 227-236, 1990.

[11] R. McKelvey and T. Palfrey, Quantal response equilibria for normal form games, 1st ed. Pasadena, Calif.: Division of the Humanities and Social Sciences, California Institute of Technology, 1994.

[12] J. Smith and G. Price, "The Logic of Animal Conflict", Nature, vol. 246, no. 5427, pp. 15-18, 1973.

[13] J. Hofbauer, K. Sigmund, "Evolutionary Games and Population Dynamics", Cambridge University Press, 1998.

[14] J. Weibull, " Evolutionary Game Theory". MIT Press, Cambridge, 1995.

[15] J. Hofbauer and W. H. Sandholm, "Stable games and their dynamics", Journal of Economic Theory, vol. 144, no. 4, pp. 1665-1693, 2009.

[16] P. Coucheney, B. Gaujal and P. Mertikopoulos, "Penalty-Regulated Dynamics and Robust Learning Procedures in Games", Mathematics of Operations Research, vol. 40, no. 3, pp. 611-633, 2015.

[17] P. Mertikopoulos and W. Sandholm, "Learning in Games via Reinforcement and Regularization", Mathematics of Operations Research, vol. 41, no. 4, pp. 1297-1324, 2016.

[18] R. T. Rockafellar and R. J.-B. Wets, Variational Analysis. Berlin: Springer-Verlag, 1998.

[19] H. Bauschke and P. Combettes, Convex analysis and monotone operator theory in Hilbert spaces, 1st ed. New York: Springer, 2011.

[20] F. Facchinei and J.-S. Pang, Finite-dimensional Variational Inequalities and Complementarity Problems. Vol. I, Springer Series in Operations Research, Springer-Verlag, New York, 2003.

[21] J. Peypouquet. Convex optimization in normed spaces: theory, methods and examples. Springer, 2015.

[22] J. B. Hiriart-Urruty and C. Lemaréchal: Fundamentals of Convex Analysis. SpringerVerlag, Berlin 2001.

[23] J. Baillon and G. Haddad, "Quelques propriétés des opérateurs angle-bornés etn-cycliquement monotones", Israel Journal of Mathematics, vol. 26, no. 2, pp. 137-150, 1977.

[24] M. Chiang, "Geometric programming for communication systems", Foundations and Trends of Communications and Information Theory, vol. 2, no. 1-2, Aug. 2005.

[25] N. Daw, J. O'Doherty, P. Dayan, B. Seymour and R. Dolan, "Cortical substrates for exploratory decisions in humans", Nature, vol. 441, no. 7095, pp. 876-879, 2006.

[26] D. Lee, "Neuroeconomics: Best to go with what you know?", Nature, vol. 441, no. 7095, pp. 822-823, 2006.

[27] J. D. Cohen, S. M. McClure, and A. J. Yu, "Should I stay or should I go? How the human brain manages the trade-off between exploitation and exploration", Philosph. Trans. Roy. Soc. B: Bio. Sci., vol. 362, no. 1481, pp. 933-942, 2007.

[28] D. Koulouriotis and A. Xanthopoulos, "Reinforcement learning and evolutionary algorithms for non-stationary multi-armed bandit problems", Applied Mathematics and Computation, vol. 196, no. 2, pp. 913-922, 2008.

[29] R. Zunino, P. Gastaldo, "Analog implementation of the softmax function", In IEEE International Symposium on Circuits and Systems, vol 2, pp II-117, 2002.

[30] A. L. Yuille and D. Geiger, "Winner-Take-All Mechanisms", In The Handbook of Brain Theory and Neural Networks, Ed. M. Arbib, MIT Press, 1995.

[31] I. M. Elfadel and J. L. Wyatt Jr., "The softmax nonlinearity: Derivation using statistical mechanics and useful properties as a multiterminal analog circuit element", In Advances in Neural Information Processing Systems 6, J. Cowan, G. Tesauro, and C. L. Giles, Eds. San Mateo, CA: Morgan Kaufmann, 1994, pp. 882-887.

[32] I. M. Elfadel, "Convex Potentials and their Conjugates in Analog Mean-Field Optimization", Neural Computation, vol. 7, no. 5, pp. 1079-1104, 1995.

[33] T. Genewein and D. A. Braun, "Bio-inspired feedback-circuit implementation of discrete, free energy optimizing, winner-take-all computations", Biological, vol. 110, no. 2, pp. 135-150, Jun. 2016.

[34] T. Michalis, "One-vs-each approximation to softmax for scalable estimation of probabilities", In Advances in Neural Information Processing Systems 29, pp. 4161-4169. 2016.

[35] P. Reverdy and N. Leonard, "Parameter Estimation in Softmax Decision-Making Models With Linear Objective Functions", IEEE Transactions on Automation Science and Engineering, vol. 13, no. 1, pp. 54-67, 2016.

[36] M. Kaisers, K. Tuyls, F. Thuijsman, S. Parsons, "An evolutionary model of multi-agent learning with a varying exploration rate (Short Paper)", Proc. of 8th Int. Conf. on Autonomous Agents and Multiagent Systems (AA-MAS 2009), Decker, Sichman, Sierra and Castelfranchi (eds.), Budapest, Hungary, pp. 1255-1256., 2009.

[37] A. Martins and R. F. Astudillo. "From softmax to sparsemax: A sparse model of attention and multi-label classification", arXiv:1602.02068 [cs.CL], Feb. 2016.

[38] D. Leslie and E. Collins, "Individual Q-Learning in Normal Form Games", SIAM Journal on Control and Optimization, vol. 44, no. 2, pp. 495-514, 2005.

[39] W. Sandholm, E. Dokumacı and R. Lahkar, "The projection dynamic and the replicator dynamic", Games and Economic Behavior, vol. 64, no. 2, pp. 666-683, 2008.

[40] R. Laraki and P. Mertikopoulos, "Higher order game dynamics", Journal of Economic Theory, vol. 148, no. 6, pp. 2666-2695, 2013.

[41] F. Alvarez, J. Bolte and O. Brahic, "Hessian Riemannian Gradient Flows in Convex Programming", SIAM Journal on Control and Optimization, vol. 43, no. 2, pp. 477-501, 2004.

[42] A. Beck and M. Teboulle, "Mirror descent and nonlinear projected subgradient methods for convex optimization", Operations Research Letters, vol. 31, no. 3, pp. 167-175, 2003.

[43] S. Shalev-Shwartz, "Online Learning and Online Convex Optimization", Foundations and Trends in Machine Learning, vol. 4, no. 2, pp. 107-194, 2011.

[44] E. Hazan, "Introduction to Online Convex Optimization", Foundations and Trends in Optimization, vol. 2, no. 3-4, pp. 157-325, 2016.

[45] A. Rangarajan, "Self-annealing and self-annihilation: unifying deterministic annealing and relaxation labeling", Pattern Recognition, vol. 33, no. 4, pp. 635-649, 2000.

[46] C. Shen and H. Li, "On the dual formulation of boosting algorithms", IEEE Trans. Pattern Anal. Mach. Intell. Feb. 25, 2010, 10.1109/TPAMI.2010.47.

[47] M. Harper, "The replicator equation as an inference dynamic", arXiv:0911.1763 [math.DS], May. 2010.

[48] Y. Sato and J. Crutchfield, "Coupled Replicator Equations for the dynamics of learning in multiagent systems", Physical Review E, vol. 67, no. 1, 2003.

[49] H. K. Khalil, Nonlinear Systems, 3rd ed., Upper Siddle River, NJ: Prentice-Hall, 2002.

[50] A. Kianercy and A. Galstyan, "Dynamics of Boltzmann Q-learning in two-player two-action games", Phys. Rev. E, vol. 85, no. 4, pp. 1145-1154, 2012.

[51] J. A. Kelner, Y. T. Lee, L. Orecchia, and A. Sidford, "An almost-linear-time algorithm for approximate max flow in undirected graphs, and its multicommodity generalizations, arXiv:1304.2338v2 [cs.Ds], Sep. 2013.

[52] R. Cominetti, E. Melo and S. Sorin, "A payoff-based learning procedure and its application to traffic games", Games and Economic Behavior, vol. 70, no. 1, pp. 71-83, 2010.

[53] K. Tuyls, K. Verbeeck and T.Lenaerts, "A selection-mutation model for Q-learning in multi-agent systems", in Proc. of the 2nd Int. Joint Conf. on Autonomous Agents and Multi-Agent Systems (AAMAS), pp. 693-700, 2003.